

Large Scale Electronic Structure Calculations with Multigrid Acceleration

E. L. Briggs, D. J. Sullivan, and J. Bernholc

Department of Physics, North Carolina State University, Raleigh, North Carolina 27695-8202

We have developed a set of techniques for performing large scale *ab initio* calculations using multigrid accelerations and a real-space grid as a basis. The multigrid methods permit efficient calculations on ill-conditioned systems with long length scales or high energy cutoffs. The technique has been applied to systems containing up to 100 atoms, including a highly elongated diamond cell, an isolated C₆₀ molecule, and a 32-atom cell of GaN with the Ga d-states in valence. The method is well suited for implementation on both vector and massively parallel architectures.

Over the course of the last several decades, plane-wave-based methods have been used to perform electronic structure calculations on a wide range of physical systems. The Car-Parrinello (CP) and other iterative methods^{1,2} have made such calculations possible for systems containing several hundred atoms.³ While these methods have been very successful, several difficulties arise when they are extended to physical systems with large length scales or containing first-row or transition-metal atoms. Special techniques have been developed to handle some of these problems. Optimized pseudopotentials,⁴⁻⁷ the augmented-wave method,⁸ plane waves in adaptive-coordinates,⁹ and preconditioning combined with conjugate-gradient techniques^{10,11} have had considerable success. However, these techniques are still constrained by the plane-wave basis set, which requires periodic boundary conditions for every system and fast Fourier transforms (FFTs) to efficiently transform between real and reciprocal space. FFTs involve non-local operations that impose constraints on the adaptability of these algorithms to massively parallel computer architectures, because they perform best on problems that can be divided into localized domains.

It has been appreciated for some time that there are potential advantages to performing electronic-structure calculations entirely in real space. Boundary conditions are not constrained to be periodic, which permits the use of non-periodic boundary conditions for clusters and a combination of periodic and non-periodic boundary conditions for surfaces. By employing nonuniform real-space grids, it is possible to add resolution locally; *e.g.*, for a surface or cluster calculation, a basis that uses a high density of grid points in the regions where the ions are located and a lower density of points in the vacuum regions can lead to order of magnitude savings in the basis

size and total computational effort.¹² More importantly, the use of a real space basis opens up the possibility of using multigrid iterative techniques to obtain solutions of the Kohn-Sham equations. Multigrid methods¹³ provide automatic preconditioning on all length scales, which can greatly reduce the number of iterations needed to converge the electronic wavefunctions. Furthermore, the real-space multigrid formulation does not involve long-range operations and is particularly suitable for parallelization and wavefunction-based $O(N)$ algorithms,¹⁴ because every operation can be partitioned into hierarchical real-space domains. In this Communication we describe the essential elements of the uniform-grid, real-space multigrid implementation as well as tests on a variety of systems, including a vacancy in diamond, a strongly elongated 96-atom diamond supercell, and a 32-atom supercell of GaN with 3d electrons in valence.

There have been a number of previous real-space grid-based electronic-structure calculations. The finite-element method was applied by White *et al.*¹⁵ to one-electron systems. They used both conjugate-gradient and multigrid acceleration to find the ground-state wavefunction. Two of the present authors¹² used nonuniform grids with locally enhanced regions in conjunction with multigrid acceleration to calculate the electronic properties of atomic and diatomic systems with nearly singular pseudopotentials. They verified that the preconditioning afforded by multigrid was effective in multi-length-scale systems. Recently, Chelikowsky *et al.*¹⁶ have used high-order finite-difference methods and soft non-local pseudopotentials on uniform grids to calculate the electronic structure, geometry, and short-time dynamics of small Si clusters.

Several issues that are absent from plane-wave or orbital-based methods arise when using a real-space grid approach. In the former case the wavefunctions, potentials, and the electronic density are representable in explicit basis functions, and thus are known everywhere. Errors in the representation of these quantities are mainly due to the truncation of the basis. In a real-space grid implementation, these quantities are known only at a discrete set of grid points, which can introduce a spurious dependence of the Kohn-Sham eigenvalues, the total energy, and the ionic forces on the positions of the ions with respect to the real-space grid. We have developed a set of techniques that can overcome these difficulties and can be used to compute accurate static and dynamical properties of large physical systems, while taking advantage of the rapid convergence rates that are made possible by multigrid methods. In our formalism the wave-

functions, density, and potentials are directly represented on a uniform three-dimensional real-space grid with linear spacing h_{grid} and number of mesh points N_{grid} . The ions are represented by soft-core norm-conserving pseudopotentials.¹⁷ Exchange and correlation effects are treated using the local density approximation (LDA) of density functional theory.

We start with the eigenvalue expression for the total electronic energy of a system of electrons and ions in the local density approximation (LDA):

$$E_{LDA} = \sum_n \epsilon_n + \int d\vec{r} \rho(\vec{r}) (\epsilon_{XC}(\rho(\vec{r})) - \mu_{XC}(\rho(\vec{r}))) - \frac{1}{2} \int d\vec{r} \rho(\vec{r}) v_{Hartree}(\vec{r}) + E_{ion-ion}. \quad (1)$$

Since the kinetic energy operator is not explicitly represented in Eq. (1), we can discretize the Kohn-Sham equations in a generalized eigenvalue formulation:

$$A[\psi_n] + B[V_{eff}\psi_n] = \epsilon_n B[\psi_n]. \quad (2)$$

A and B are the components of the *Mehrstellen* discretization,¹⁸ which is based on Hermite's generalization of Taylor's theorem. It uses a weighted sum of the wavefunction and potential values to improve accuracy of the discretization of the *entire* differential equation, not just the kinetic energy operator. The weights are listed in Table I. The Kohn-Sham hamiltonian, wavefunctions, and eigenvalues are represented to $O(h_{grid}^4)$, but the prefactor for the $O(h_{grid}^4)$ term is smaller than 0.01. Alternative grid-based discretizations of the Kohn-Sham equations do exist, *e.g.*, high-order central-finite-difference methods were used by Chelikowsky *et al.*¹⁶ They discretize the kinetic energy operator only, and typically use an $O(h_{grid}^{12})$ discretization. *Mehrstellen* discretization differs from central-finite-difference discretization in that it achieves higher accuracy by using more *local* data. For example, the fourth-order A operator extends to second nearest neighbors; in contrast, the twelfth-order central-finite-difference operator extends linearly out 6 mesh points in each direction. We have found in our test cases that the fourth-order *Mehrstellen* discretization produces equivalent or better accuracy than the standard sixth-order finite-difference discretization.

When using plane-wave basis sets, the accuracy is usually determined by the convergence of the total energies and eigenvalues as the energy cutoff used is increased. In calculations employing a real-space grid, it is also necessary to consider the dependence of the energies and the ionic forces on the positions of the ions relative to the real-space grid points. This non-physical dependence is not present in plane-wave methods. Our selection of the *Mehrstellen* operator and discretization techniques is designed to reduce this grid dependence as much as possible, because the determination of accurate trajectories in molecular dynamics simulations and the calculation

of dynamical quantities is very sensitive to this dependence. Our approach is to perform calculations on isolated atoms and to monitor the variations as the ion is moved relative to the grid points. To reduce finite-size effects, periodic boundary conditions and a large supercell are used in these tests rather than cluster boundary conditions. In Fig. 1 we show the variation in the total energy of a carbon atom. For bulk calculations, we also test the grid resolution by monitoring the total energy as the atomic lattice is rigidly translated with respect to the grid.

The solution of Eq. (2) by direct matrix methods or by standard iterative techniques is prohibitively expensive for very large systems because they require $O(N_{grid}^{5/3})$ operations per wavefunction on an uniform mesh of N_{grid} points.¹⁹ As the grid resolution is increased, *e.g.*, to treat systems containing first row or transition metal atoms, the rate of convergence becomes unacceptable. A similar slow-down is present in plane-wave methods as the energy resolution increases.

To efficiently solve Eq. (2), we have used multigrid iteration techniques that accelerate convergence by employing a sequence of grids of varying resolutions. The final solution is obtained on a grid fine enough to accurately represent the pseudopotentials and the electronic wavefunctions. If the solution error is expanded in a Fourier series, it may be shown that iterations on any given grid level will quickly reduce the components of the error with wavelengths comparable to the grid spacing but are ineffective in reducing the components with wavelengths large relative to the grid spacing.^{13,19} The solution is to treat the lower frequency components on a sequence of auxiliary grids with progressively larger grid spacings, where errors appear as high frequency components. This procedure provides excellent preconditioning for all length scales present in a system and leads to very rapid convergence rates. The operation count to converge one wavefunction with a fixed potential is $O(N_{grid})$, compared to $O(N_{grid} \log N_{grid})$ for FFT-based approaches.² In addition, all operations except orthogonalization are short ranged, which allows for easy parallelization and a natural implementation of $O(N)$ algorithms.¹⁴

One iteration towards the solution of Eq. (1) consists of a multigrid step to solve Eq. (2), followed by orthogonalization of the orbitals, and an update of the electronic density. The components of the effective potential that depend on the density (Hartree and exchange-correlation) are then recomputed for the new density and the process is repeated until self-consistency is reached. We have found that it is often necessary to implement a mixing procedure when updating the electronic density to minimize charge sloshing.²⁰ The Hartree potential is computed by solving Poisson's equation using multigrid iterations on the corresponding *Mehrstellen* discretization. The non-local ionic pseudopotential is evaluated in real space using the Kleinman-Bylander²¹ separable form. To facilitate comparisons with plane-wave calculations, we

define an equivalent energy cutoff for the multigrid calculation, $\pi^2/2h^2$ [Ry], to be equal to that of a plane-wave calculation that uses a FFT grid with the same spacing as the multigrid calculation.²² Using this convention, the computational time required to perform one step of the above procedure is approximately the same as for a single step in the CP method when the cutoffs are equal.

The multigrid method outlined above has been tested on several large systems. A 64-atom diamond supercell was chosen as a representative periodic system for which both multigrid and Car-Parrinello calculations are feasible. The calculations were performed on the perfect crystal and the relaxed vacancy using the same pseudopotential.¹⁷ The CP cutoff was 35 Ry, while the multigrid code used a grid spacing of 0.336 bohr, which results in a grid with the same total number of points as the FFT grid needed in the CP code. The k-space sampling was restricted to the Γ point. The dynamical relaxation method²³ was used to quickly relax the vacancy in both calculations.

The Car-Parrinello and real-space calculations are in excellent agreement with each other. The relaxed ionic positions agree to within 0.009 bohr for all ions in the supercell, and the largest difference in Kohn-Sham eigenvalue is 0.06 eV. The diamond vacancy introduces an initially triply degenerate level in the gap, which splits into a singlet and a doublet due to the Jahn-Teller effect. The magnitude of this splitting is 0.32 eV in both the real-space and Car-Parrinello results. The main results are summarized in Table II. The discrepancies between the LDA calculations and experimental values for the bandgap and the cohesive energy are well-known deficiencies of density functional theory.

An isolated C_{60} molecule was selected as an example of a non-periodic system. The simulation cell was a cube of length 23 bohr and the grid spacing was 0.360 bohr. The initial ionic coordinates were generated using the classical Tersoff-Brenner potential,²⁴ and the electronic wavefunctions were set to zero on the boundaries of the cell. After the convergence of the electronic system, the ions were relaxed using the same relaxation scheme as before. Two distinct bond lengths were found in the final structure, corresponding to the carbon-carbon single and double bonds. There were twice as many single bonds as double bonds, and the average double and single bond lengths were 1.39 and 1.44 Å, compared to 1.41 and 1.45 Å obtained in a previous CP calculation²³ for the C_{60} solid. The standard deviations of the bond-length distributions were on the order of 10^{-3} Å in both calculations. The experimental values for the solid are 1.40 and 1.45 Å, respectively.

A significant advantage of the multigrid method is the speed of convergence to the electronic ground state for a given initial ionic configuration. For systems requiring a small energy cutoff or of small size, the speed advantage with respect to CP-based methods is not substantial. However, for systems requiring a large energy cutoff, or of large dimensions, this advantage—as measured in

actual computational time—is typically an order of magnitude. This is because the maximum stable timestep in the Car-Parrinello method must be much smaller than in the multigrid approach.

To illustrate the ability of the multigrid method to handle ill-conditioned systems, test calculations were performed on a 32-atom GaN cell in the zinc-blende structure with the Ga d-electrons in valence, and on a highly elongated 96-atom diamond cell of dimensions (6.72, 6.72, 80.64) bohr. For the GaN cell a grid spacing of 0.175 bohr was used, which corresponds to an energy cutoff of 160 Ry in a plane-wave calculation. Only the Γ point was included in k-space sampling. Starting with random initial wavefunctions, 240 self-consistent steps were required to converge the total energy. Recently, several calculations have been performed on GaN that explicitly included the d-electrons in valence; the multigrid results are in good agreement with these calculations (see Table III). The calculations for the elongated diamond cell used a grid resolution of 0.336 bohr, corresponding to an energy cutoff of 35 Ry in a plane wave calculation. For convenience, an approximate solution was generated on a grid with twice the final spacing, on which each iteration is eight times faster than on the final grid. After transferring this approximate solution to the final grid, only 50 self-consistent iterations were sufficient to converge the total energy to a relative tolerance of 10^{-8} .

These two ill-conditioned systems represent worst-case scenarios, and the performance of the multigrid method is considerably better for typical systems. Table IV illustrates the convergence properties as a function of grid resolution for an 8-atom diamond cell. The observed convergence rates are largely independent of the energy cutoff. The number of self-consistent steps required to converge the density is also nearly independent of the system size. At an equivalent cutoff of 35 Ry, the multigrid method required 17 steps to converge the total energy of the 8-atom diamond cell to a tolerance of 10^{-8} . When the same calculation was performed on the 64-atom cell only 20 SCF steps were needed. This favorable scaling with respect to both energy cutoff and system size, and the inherent stability of multigrid methods offer the possibility of performing electronic-structure calculations on very large systems.

In summary, we have developed a methodology for performing large-scale *ab initio* electronic-structure calculations entirely in real space. The method uses highly-efficient multigrid techniques to accelerate convergence rates, which is particularly important for ill-conditioned systems requiring high energy cutoffs or with large length scales. In addition, the multigrid method is readily adaptable to parallel computer architectures, and its block structure makes it very suitable for $O(N)$ approaches.

- ¹ R. Car and M. Parrinello, Phys. Rev. Lett. **55**, 2471 (1985).
- ² For a recent review, see M. C. Payne, D. C. Allan, M. P. Teter, T. A. Arias, and J. D. Joannopoulos, Rev. Mod. Phys. **64**, 1045 (1992).
- ³ See, for example, I. Stich, M. C. Payne, R. D. King-Smith, and J.-S. Lin, Phys. Rev. Lett. **68**, 1351 (1992); K. D. Brommer, M. Needels, B. E. Larson, and J. D. Joannopoulos, *ibid.* **68**, 1355 (1992); P. Bogusławski, Q.-M. Zhang, Z. Zhang, and J. Bernholc, *ibid.* **72**, 3694 (1994).
- ⁴ D. Vanderbilt, Phys. Rev. B **41**, 7892 (1990).
- ⁵ A. M. Rappe, K. M. Rabe, E. Kaxiras, and J. D. Joannopoulos, Phys. Rev. B **41**, 1227 (1990).
- ⁶ J.-S. Lin, A. Qteish, M. C. Payne, V. Heine, Phys. Rev. B **47**, 4174 (1993).
- ⁷ G. Li and S. Rabii (1992), unpublished.
- ⁸ P. E. Blöchl, Phys. Rev. B **41**, 5414 (1990).
- ⁹ F. Gygi, Europhys. Lett. **19**, 6617 (1992).
- ¹⁰ T. A. Arias, M. C. Payne, and J. D. Joannopoulos, Phys. Rev. Lett. **69**, 1077 (1992); Phys. Rev. B **45**, 1538 (1992).
- ¹¹ M. P. Teter, M. C. Payne, and D. C. Allan, Phys. Rev. B **40**, 12255 (1989).
- ¹² J. Bernholc, J.-Y. Yi, and D. J. Sullivan, Faraday Disc. Chem. Soc. **92**, 217 (1991).
- ¹³ A. Brandt, Math. Comp. **31**, 333 (1977); GMD Studien, **85**, 1 (1984).
- ¹⁴ G. Galli and M. Parrinello, Phys. Rev. Lett. **69**, 3547 (1992); F. Mauri, G. Galli, and R. Car, Phys. Rev. B **47**, 9973 (1993); P. Ordejón, D. A. Drabold, M. P. Grumbach, and R. M. Martin, *ibid.* **48**, 14646 (1993).
- ¹⁵ S. R. White, J. W. Wilkins, and M. P. Teter, Phys. Rev. B **39**, 5819 (1989).
- ¹⁶ J. R. Chelikowsky, N. Troullier, and Y. Saad, Phys. Rev. Lett. **72**, 1240 (1994); J. R. Chelikowsky, N. Troullier, K. Wu, and Y. Saad, Phys. Rev. B **50**, 11355 (1994); X. Jing, N. Troullier, D. Dean, N. Binggeli, J. R. Chelikowsky, K. Wu, and Y. Saad, *ibid.* **50**, 12234 (1994).
- ¹⁷ D. R. Hamann, Phys. Rev. B **40**, 2980 (1989).
- ¹⁸ L. Collatz, *The Numerical Treatment of Differential Equations*, (Springer-Verlag, Berlin, 1960), p. 164.
- ¹⁹ W. L. Briggs, *A Multigrid Tutorial*, (SIAM Books, Philadelphia, 1987).
- ²⁰ See, for example, W. E. Pickett, Comp. Phys. Rep. **9**, 115 (1989).
- ²¹ L. Kleinman and D. M. Bylander, Phys. Rev. Lett. **48**, 1425 (1982).
- ²² We assume the electronic charge density is expanded with energy cutoff twice as large as that of the wavefunctions.
- ²³ Q.-M. Zhang, J.-Y. Yi, and J. Bernholc, Phys. Rev. Lett. **66**, 2633 (1991).
- ²⁴ J. Tersoff, Phys. Rev. Lett. **56**, 632 (1986); Phys. Rev. Lett. **61**, 2879 (1988); Phys. Rev. B **37**, 6991 (1988); D. W. Brenner, Phys. Rev. B **42**, 9458 (1990).
- ²⁵ C. Kittel, *Introduction to Solid State Physics*, (John Wiley and Sons, New York, 1986), p. 55.
- ²⁶ V. Fiorentini, M. Methfessel, M. Scheffler, Phys. Rev. B **47**, 13353 (1993).
- ²⁷ A. F. Wright, and J. S. Nelson, Phys. Rev. B **50**, 2159 (1994).

FIG. 1. Variation of the total energy as a carbon atom is displaced relative to the grid point along a coordinate axis. The grid spacing h is 0.336 bohr.

TABLE I. *Mehrstellen* discretization weights in 3D (h is the grid spacing).

Grid point	$6h^2 \cdot A$	$6 \cdot B$
central	-24	6
nearest neighbors	2	1
second neighbors	1	0

TABLE II. Comparison of perfect crystal and vacancy in diamond results obtained using 64-atom supercells [eV].

	Car-Parrinello ^a	Multigrid	Experiment ^b
Perfect Crystal			
Band gap	4.53	4.53	5.50
Cohesive energy	8.49	8.54	7.37
Vacancy			
Formation energy	6.98	7.07	

^a35 Ry cutoff and Γ point sampling.

^bKittel.²⁵

TABLE III. Comparison of GaN results obtained with Ga 3d electrons in valence. See text [eV].

	Band Gap	Cohesive Energy
Fiorentini <i>et al.</i> ^a	2.00	10.89
Wright <i>et al.</i> ^b	1.89	
Multigrid ^c	1.79	10.68

^aFull potential LMTO.²⁶

^bPlane wave calculation with 240 Ry cutoff.²⁷

^c32-atom supercell, 160 Ry cutoff and Γ point sampling.

TABLE IV. Multigrid convergence tests in an 8-atom supercell: grid spacing [bohr], equivalent plane-wave cutoff [Ry], total energy [a.u.], and the number of steps to converge the density.

Grid Spacing	Cutoff	Total Energy	SCF steps
0.421	25	-44.87968	22
0.336	35	-45.03331	17
0.280	60	-45.06240	21
0.210	110	-45.06605	26




Figure 1.